

Pills to Puffs:

The Relationship Between Cannabis and Prescription Drugs

Maxwell Tabarrok, University of Virginia

maxwell.tabarrok@gmail.com

<https://www.maximumprogress.org/about-me>

Abstract

In 2001, recreational use of cannabis was illegal in every US state. By 2020, it was legal in 20 of them. Over the same period, cannabis use in America doubled. Previous research suggests that cannabis has significant substitution or complementarity effects with several other substances including alcohol ([Subbaraman, 2016](#)), tobacco ([Miller, Seo, 2021](#)), other illicit drugs ([Powell, Pacula, Jacobson, 2018](#)), and prescription medicines ([Raman, Bradford, 2022](#)). In this paper I extend the literature on the interactions between cannabis and prescription drugs in 3 ways. First, I use synthetic controls to study the effect of legalization on both cannabis and prescription medicine consumption. Second, I use the event study specification in [Callaway and Sant'Anna \(2021\)](#) to improve upon sometimes arbitrary matching procedures in two way fixed effects models. Finally, I have hosted [an interactive appendix where you can reproduce all of the results](#) for all combinations of states, controls, and prescription drug categories. I find that legalization increased cannabis consumption in legal states by about 30% on average. Only anxiety and depression medications see statistically significant decreases in response, although all drug classes have a negative point estimate. The substitution effect on prescription medicine consumption is milder than suggested by previous research.

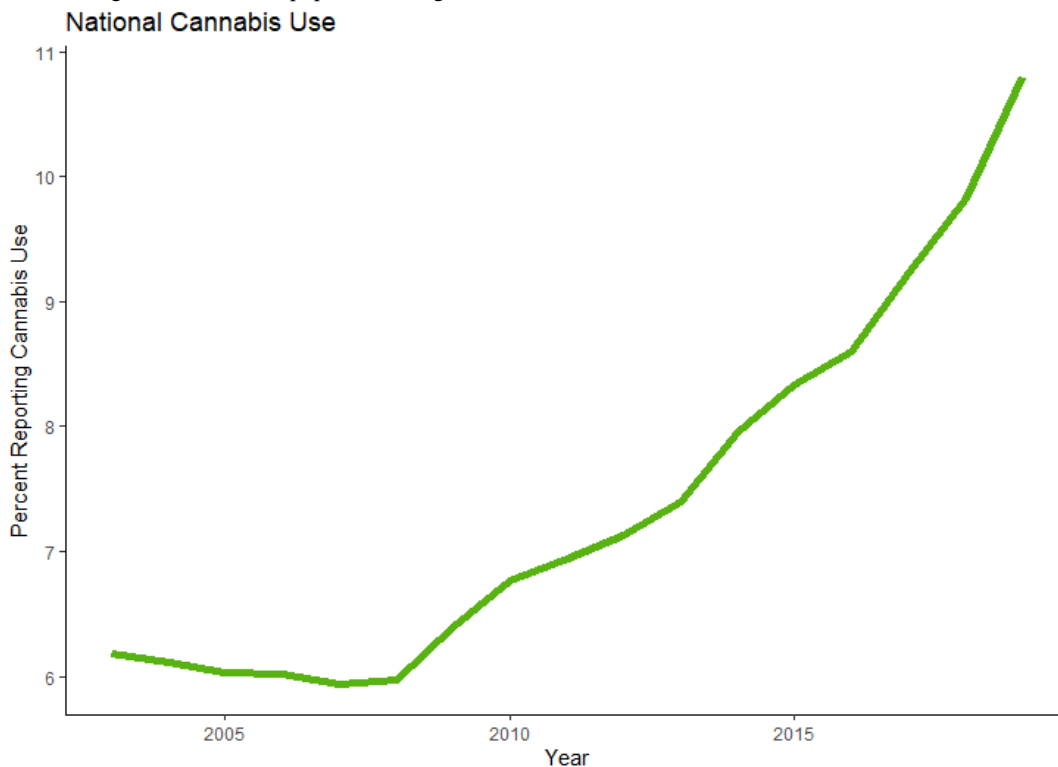
Introduction

Cannabis use in the United States has increased significantly in the last 20 years as an already enormous black-market industry grew with new legal status. Cannabis use has strong statistical relationships with other substance use, including prescription medicines. According to the

Figure 1: National Average Proportion of Respondents Reporting Cannabis Use

Source: [National Survey of Drug Use and Health \(NSDUH\)](#)

Note: Average over states, not population weighted.



[National Survey of Drug Use and Health \(NSDUH\)](#), people who report using cannabis at least once in the past year are 50% more likely to report taking prescription medicines for mental health conditions. This is obviously not a pure causal effect, but it suggests that cannabis' relationship with prescription drugs is important for understanding the effects of legalizing cannabis use.

The causal relationship between cannabis and prescription medicines is mediated through complementarity and substitution effects. If hotdogs were a Schedule I controlled substance, we would see fewer hotdog buns on the shelves and more burgers. So it is with cannabis. As cannabis is legalized in more states, the cost of consumption decreases as users no longer have to face legal risk to consume the drug. Quality control may also increase as brands establish public reputations. An increase in the demand for cannabis will increase the use of complements and reduce the use of substitutes.

Previous research suggests that cannabis serves as a substitute for a broad range of prescription drugs, especially pain relievers and mental health drugs. A survey of Canadian medical cannabis users found that more than half of them (63%) reported using cannabis as a substitute for a prescription medication ([Lucas, 2017](#)). Several papers have studied cannabis as a substitute for pain medications in the context of the ongoing opioid crisis. [Bachhuber \(2014\)](#) and [Powell \(2018\)](#) both find significant decreases in opioid related deaths following the passage of medical marijuana laws. [Raman, Bradford \(2022\)](#) study a wider range of pharmaceutical products using a difference-in-difference following recreational cannabis laws and find that cannabis has substitution effects with pain medicines but also for medicines prescribed for anxiety, depression, sleep, seizures, and psychosis.

I use several methods to extend and verify the previous research on this question. First, I use difference-in-difference, synthetic controls, and an event study to show that cannabis legalization caused a significant increase in cannabis consumption, verifying substitution as a mechanism for changes in other consumption habits. Then, the same methods allow me to use cannabis legalization as a quasi-experimental test of the elasticity of substitution between cannabis and prescription medicines. These new data and methods temper results from previous research.

Did Legalization Increase Cannabis Use?

Before using legalization to test the elasticity of substitution between cannabis and prescription drugs, I have to test the relationship between legalization and cannabis use. Intuition and economic theory both predict that legalization will increase use. Criminalizing a product can be modeled as a type of tax, so an increase in consumption after removing that tax will not be surprising. Still, it is worth testing the causal relationship empirically to get a sense of the magnitude of this effect. Will legalization triple cannabis use? Or just increase it modestly?

The data on cannabis use come from the National Survey on Drug Use and Health (NSDUH). This is a yearly, nationally representative survey which measures drug use habits. The metric I track is the percentage of respondents reporting marijuana use in the past month. The use of survey data to measure cannabis consumption is a necessary compromise for studying the effect of legalization. Compared to other measures of consumption, like sales or tax data, surveys have extra noise and possible bias as legalization may change the willingness of respondents to report cannabis use as well as actual consumption rates. Unfortunately, these other measures are not available before legalization so they can't be used to estimate its effect.

Cannabis use certainly increased in the 9 states I study which legalized before 2017. However, cannabis use is growing in all states, not just the ones that legalized. And all the legal states had growing consumption even before they legalized. Table 1 below shows that cannabis consumption in all legal states grew about 50% which is higher than the national average. This is consistent with legalization causing more consumption, but it isn't enough to prove a causal effect or measure its magnitude.

Average Proportion of Population Reporting Cannabis Use

Before and After Legalization


Pre-Legal Means	Post-Legal Means	State	Change
7.99%	14.55%	Washington	+6.56%
9.09%	16.28%	Colorado	+7.19%
9.67%	17.41%	Oregon	+7.73%
10.61%	16.19%	Alaska	+5.57%
7.98%	12.81%	California	+4.83%
9.81%	16.76%	Maine	+6.95%
7.18%	14.93%	Nevada	+7.75%
9.32%	14.67%	Massachusetts	+5.36%
10.20%	17.11%	District of Columbia	+6.91%

Pairwise DiD

One of the simplest methods to get an estimate of the magnitude of the causal effect of legalization is difference-in-difference (DiD). I noted above that comparing the pre and post legalization averages is not enough to recover the causal impact of legalization for two reasons: Other non-legal states also increased consumption over the same period, and legal states before legalization increased consumption. To get a better estimate of the causal effect of legalization, we need to consider the fact that some growth in cannabis consumption would be expected even with no legislative changes. DiD combines cross-state and cross-time comparisons to subtract the increase in consumption expected from these sources, leaving only the causal effect of legalization.

Below is an example of a simple DiD.

Figure 2: Colorado and Idaho Cannabis Use TWFE



This is a comparison between cannabis consumption in Colorado (pink) and Idaho (gray). We can see that before Colorado legalized cannabis in 2012, consumption in Colorado and Idaho moved in parallel, always separated by about a 5% difference. After 2012, Colorado’s cannabis consumption increased quickly from around 10% to more than 16%. Idaho’s consumption also increased but not by nearly as much. If Colorado and Idaho had stayed in their parallel trends, we would have expected Colorado to end up at about 11% consumption, 5 percentage points above Idaho, which is shown by the blue line. Comparing post-legalization Colorado to the counterfactual line removes the increase in cannabis consumption that we already expected to happen based on the trends in a control state. Then any increase beyond that line is attributed to something that changed in 2012: legalization.

Below are the regression results for the DiD comparison between Colorado and Idaho that I graphed above. The left-hand side variable in this regression equation is cannabis use. The right-hand side is just two dummy variables and their interaction. The variable “treated” is 1 when an observation comes from Colorado, which is the state that gets “treated” by legalization. “Time” is 1 when an observation comes on or after 2012, when Colorado’s treatment began. Therefore, “treated” contains the level differences in cannabis consumption between Colorado and Idaho in

Table 1: Colorado and Idaho Cannabis Use TWFE

	Dependent variable: <i>Cannabis Use</i>			
	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	0.056900	0.004453	12.778	4.12e-14***
treated	0.033995	0.006298	5.398	6.24e-06***
time	0.017155	0.006680	2.568	0.0151*
treated:time	0.054712	0.009446	5.792	1.98e-06***

*Note: Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1*
Residual standard error: 0.01408 on 32 degrees of freedom
Multiple R-squared: 0.8963, **Adjusted R-squared:** 0.8866
F-statistic: 92.23 on 3 and 32 DF, **p-value:** 7.804e-16

the pre-treatment period, and “time” measures the increase in cannabis consumption observed in Idaho after 2012.

Counterfactual Colorado is exactly the sum of “Intercept,” “time,” and “treated.” This is how Colorado’s cannabis consumption would have changed without legalization, taking into account their higher level relative to Idaho and the general upward trend in consumption experienced everywhere. Then the interaction term is the difference between how much we expected Colorado’s consumption to grow over time and how much it actually grew. This effect, about a 5 percentage point increase in consumption, is attributed to the causal effect of Colorado legalizing cannabis in 2012.

The reader can compare all the legal states I study to any choice of control with this graph and regression output at [this online appendix which reproduces all of my results](#). All the legal states, when compared to appropriate controls, have results similar to Colorado. Although the graphs can be convincing, this verification process feels a bit ad hoc. How do we decide which states to compare? How parallel is parallel enough? How do we aggregate the pairwise results into one average metric for the causal effect of cannabis legalization?

These questions will lead us to the next two quasi-experimental methods that I use to analyze this question.

Synthetic Controls

The synthetic control method is a way of formalizing the choice of comparison state in DiD. When you’re flipping through the options for control states in a pairwise DiD you’re looking for a state that matches the legal state’s trend in the pre-treatment period well. Visually inspecting the graph is enough to find a pretty good match, but sometimes none of the options are great or there’s a large set of “pretty good” controls but it’s not clear which is the best one. Synthetic controls addresses both of these problems.

The basic idea is to construct the ideal control state by finding the weighted average of all the control states which best matches the treatment state’s trend in the outcome variable before it gets treated. The pairwise DiDs we looked at above are just edge cases of this process where the weights for all the controls are set to zero except one. For the technical details of how this optimization works see [Abadie \(2021\)](#) and for an [applied walkthrough in R go here](#).

In our context this means finding the weighted average of states that have not legalized cannabis in 2019 which best matches a legal state’s trend in cannabis consumption in the several years before they legalized. Thanks to my computer’s linear algebra skills this is not too difficult. I get results that closely match the direction and magnitude I found with pairwise DiDs.

Most states see a noticeable bump in cannabis consumption of about 4 percentage points compared to their synthetic controls after they legalize.

Figure 3: Colorado Cannabis Use Synthetic Control



A few, like California, see a smaller bump and Alaska’s cannabis consumption doesn’t seem to

separate much from its synthetic version. You can look through [the synthetic control results for all of the states on the online appendix I hosted here](#).

Figure 4: California Cannabis Use Synthetic Control


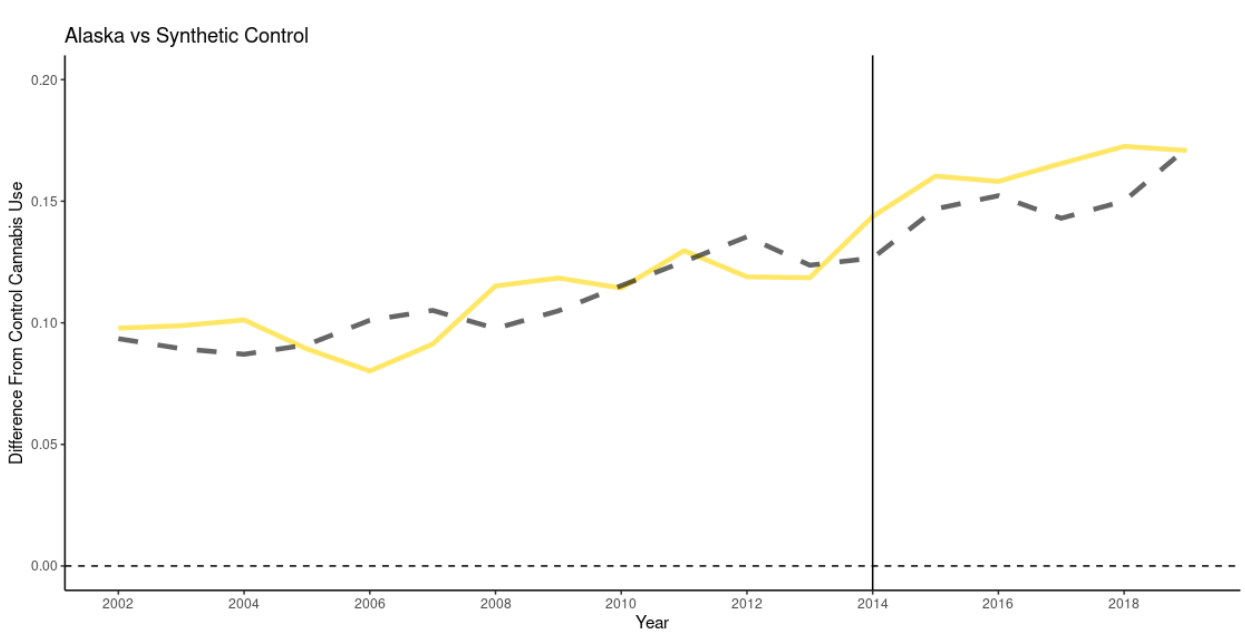


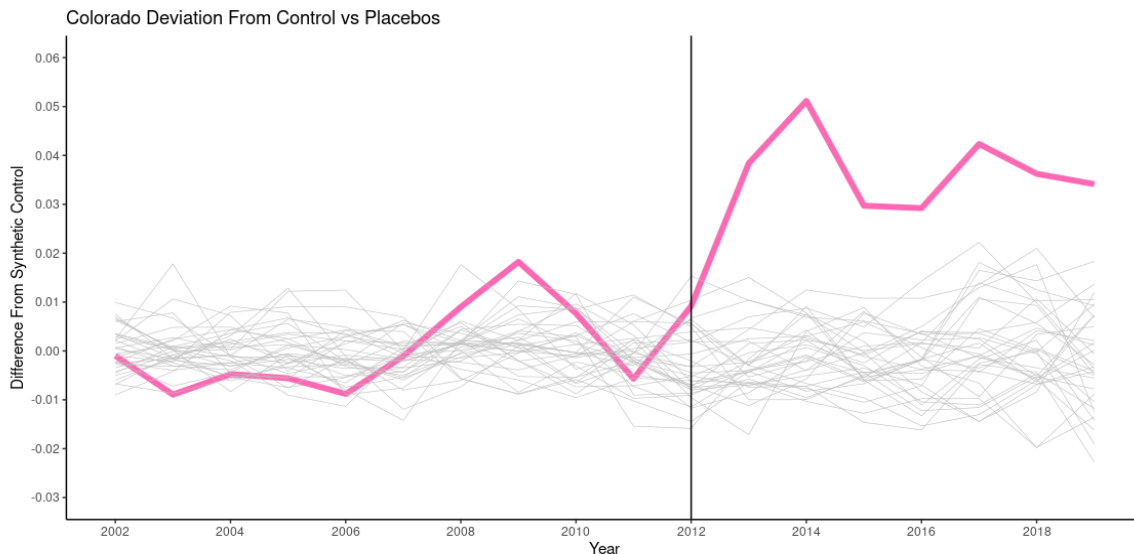
Figure 5: Alaska Cannabis Use Synthetic Control



Statistically minded readers might be wondering whether we can conclude anything from these results. After all, we only optimized the synthetic control to match the treatment state for the years before legalization. We should reasonably expect the two states to start to diverge more when they leave the optimization window. So how big of a divergence do we need to conclude that we're really observing an effect of a policy change, and not just random variation?

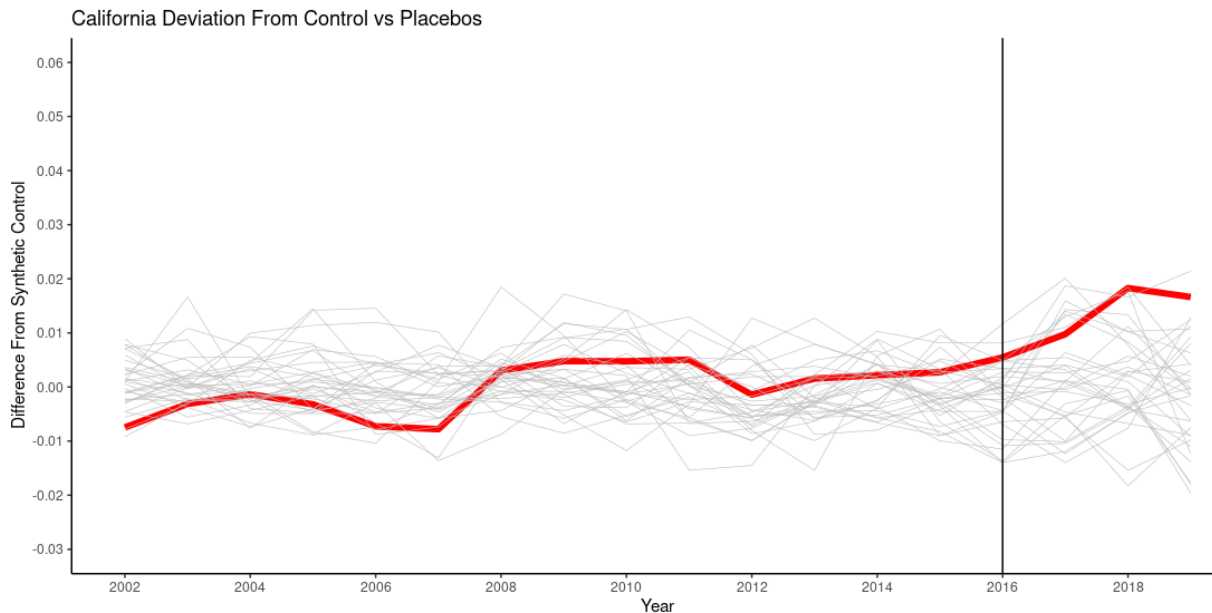
Synthetic controls answers this question using placebo tests. We'll repeat the same process as above, optimizing over control states to match the state-of-interest as closely as possible in the pre-treatment period, but apply it to states which never legalized cannabis. After applying the synthetic control process to all the control states, I will plot the differences between the placebo's paths and the path of their synthetic versions. Any divergence between them is caused by random variation, not by relevant legislative changes. If we compare the range of differences that result from random variation to the same difference for the legal state we can get a sense of how unlikely the observed divergence in the legal state would be if legalizing cannabis had no impact. For Colorado we'd get a graph that looks like this:

Figure 6: Colorado Cannabis Use Compared to Placebos



The pink line on this graph is the difference between the two lines on the first synthetic control plot I showed for Colorado. Each of the gray lines is the same difference but for a state which never legalized cannabis. We see that even the never-legal states do diverge somewhat from their synthetic controls, but Colorado's divergence is clearly separated from this pack. As expected, California doesn't see as dramatic of a bump. But it does rise to the top of the cluster of lines which, given our theoretical background knowledge of the effects of legalization, is suggestive of a positive effect.

Figure 7: California Cannabis Use Compared to Placebos



The placebo tests for all the states can be viewed on [the online appendix](#). Flipping through the results you'll see that all the states have positive results and most of them clearly separate from the random variation lines. That gives some confidence that the overall impact of legalization on cannabis is solidly positive but what exactly is the size of the effect? How do we aggregate these state-by-state results across all the states we observe? These questions are answered by the next method I use.

Event Study

Event studies, or dynamic DiDs, are extensions on the simple DiD model we used in the first section. The simple model has only three variables: one dummy variable tracking which state gets the treatment, another tracking when it happens, and their interaction. An event study DiD includes fixed effects for all the states and time periods along with an interaction term which is one when the current state is x periods away from the treatment start ([Sun, Abraham, 2021](#)). By estimating this equation for various values of x , you can estimate treatment effects for any year around the treatment for which you have data.

This specification addresses two limitations of the models we looked at above. First, we can include all of the legal states and all available control states in a single regression. The relevant time variable in the event study specification is not 'year' but 'years since treatment'. So even states with different treatment times can be aligned and their effect sizes averaged into a single estimate. Second, we can account for delayed effects of a treatment. In the simple DiD we just get a single average of cannabis consumption over the entire post-legalization period. But if it took a few years for cannabis consumption to increase after legalization, that average would be brought down by those first few years and we might not see a significant treatment effect even if there is a large one that is just delayed. Since the rollout of commercial cannabis sales is often slow following official legalization this might be relevant for our context.


However, recent work by [Callaway and Sant'Anna \(2021\)](#) has highlighted potential biases in the standard event study model. Specifically, the model may suffer from "pivotal unit" bias, where the estimated treatment effect is driven by a small number of influential units that are disproportionately affected by the treatment. This bias can arise when the treatment effect varies

across units or when there are important time-varying confounders that are not fully controlled by the fixed effects.

To address this issue, Callaway and Sant'Anna propose a modified event study model that uses a data-driven approach to identify pivotal units and downweight their influence on the estimated treatment effect ([they also provide a convenient R package](#)). The key innovation of their method is to use a lasso penalty to select a subset of units that are most likely to be pivotal, and then to estimate the treatment effect using only these units. The lasso penalty encourages sparsity in the selection of pivotal units, which helps to avoid overfitting and reduce the risk of selecting irrelevant units.

When I apply the modified event study model I get this result

Figure 8: Average Treatment Effect of Legalization on Cannabis Use



Each of the dots in this graph is an average across all the treated states of the coefficients on the interaction term which tracks whether state i is x years away from treatment. Before the treatment starts these coefficients are all close to zero which is good. Simply being a state that

eventually legalizes should not have an impact on cannabis consumption before legalization actually happens.

After legalization, cannabis use immediately jumps but it doesn't seem to get to its new equilibrium immediately. It takes at least 3 years for cannabis use to level out at 3 percentage points above non-legal states. The graph shows that cannabis use may still be increasing 7 years into the treatment, although only Colorado and Washington have that much data. The overall average treatment effect across all the states, reported in the top row of the table, is almost

Table 2: Average Treatment Effect of Legalization on Cannabis Use

Dependent variable: <i>Cannabis Use</i>				
	Estimate	Std. Error	[95% Conf. Int.]	Significance
ATT	0.0296	0.003	0.0237–0.0355	*
Dynamic Effects:				
<i>Time</i>	Estimate	Std. Error	[95% Simult. Conf. Band]	Significance
-16	0.0024	0.0009	0.0001–0.0047	*
-15	-0.0020	0.0024	-0.0084–0.0045	
-14	-0.0017	0.0018	-0.0063–0.0030	
-13	0.0013	0.0031	-0.0070–0.0096	
-12	0.0004	0.0024	-0.0059–0.0068	
-11	-0.0007	0.0030	-0.0086–0.0073	
-10	0.0002	0.0015	-0.0039–0.0042	
-9	-0.0015	0.0022	-0.0074–0.0043	
-8	0.0058	0.0022	0.0000–0.0117	*
-7	0.0032	0.0019	-0.0018–0.0082	
-6	-0.0006	0.0031	-0.0089–0.0078	
-5	0.0035	0.0018	-0.0013–0.0083	
-4	0.0018	0.0021	-0.0037–0.0073	
-3	0.0080	0.0029	0.0004–0.0156	*
-2	0.0014	0.0023	-0.0047–0.0075	
-1	0.0009	0.0013	-0.0026–0.0043	
0	0.0103	0.0035	0.0011–0.0196	*
1	0.0204	0.0056	0.0056–0.0353	*
2	0.0292	0.0074	0.0096–0.0488	*
3	0.0266	0.0068	0.0086–0.0446	*
4	0.0310	0.0050	0.0177–0.0444	*
5	0.0324	0.0043	0.0210–0.0438	*
6	0.0411	0.0028	0.0338–0.0483	*
7	0.0457	0.0024	0.0395–0.0520	*

exactly 3%. Among the legal states their cannabis consumption prevalence was less than 10% on average before legalization so this 3 percentage point increase corresponds to more than a 30% increase in consumption prevalence. You can verify this directly by running this same model in logs in [the online appendix](#).

All the tests we saw agree in direction and magnitude and they also fit our expectations from intuition and economic theory. We can confidently conclude that cannabis legalization has a large positive effect on the prevalence of cannabis use. With this knowledge we can now use legalization to test the impact of cannabis use on prescription medicine consumption. We know that legalization will cause an increase in use, but it doesn't come with the confounding selection effects that simply looking at people who use more cannabis has.

Does Cannabis Substitute For Prescription Drugs?

Data Description

Answering this question requires first collecting a measure of prescription drug use. This comes with a few challenges. First, America's decentralized health system makes collecting comprehensive data difficult. Medicaid, on the other hand, collects detailed price and volume data for all of the prescriptions it fills through the [State Drug Utilization Data \(SDUD\) archive](#). Although users of this program are not a random sample of prescription drug users, Medicaid comprises a large enough fraction of US health spending ([almost 20%](#)) that even effects which don't generalize to other users of health insurance are large enough to be important.

The second challenge is deciding which drugs to study. There is some evidence that cannabis has positive effects on a wide range of conditions including depression, ([Feingold, Weinstein, 2021](#))

anxiety ([Blessing, et al, 2015](#)), pain, and seizures ([Whiting, et al, 2015](#)). But for most conditions that prescription medicines are meant to address, like bacterial infections, cancer, or viruses, there is no evidence that cannabis has any effect. Looking at overall prescription drug use would wash out the possible substitution effects in areas where cannabis might replace prescription drugs with lots of noise from antibiotics and cancer drugs that have no relationship with cannabis use. The broad drug classes that are most often referenced in relation to cannabis and the ones I will study are depression, anxiety, sleep, seizures, pain, spasticity, psychosis, and nausea.

For the selection of relevant classes of drugs and the sorting of drugs into those classes I follow David Bradford ([Bradford, 2017](#)), Shyam Raman, and Ashley Bradford ([Raman, Bradford, 2022](#)) in their papers on cannabis laws who kindly provided their data on these drug classes. Their starting point is the clinical classes listed in drug's FDA approvals, also known as their 'on-label' indications. However, they point out that a large fraction of prescription medicines are prescribed 'off-label' by physicians. For example, only looking at drugs that are 'on-label' for anxiety would exclude depression medications and beta-blockers which are often prescribed to treat symptoms of anxiety. The appendices of the paper's references above detail their process for tracking both on-label and the most common off-label drugs for each of the 8 drug classes that I study.

The final challenge is that the absolute figures for Medicaid prescriptions are highly skewed by legislative changes which expanded the role of Medicaid, especially the introduction of Part D in 2006. To control for this, I also collect data on the [number of Medicaid enrollees for each state from MACPAC](#) and use per-enrollee metrics for all of the outcome variables. The main outcome variable we will be plotting in all the following graphs is the number of units (usually pills)

reimbursed in each state and drug class per Medicaid enrollee in that state for every year from 2008 to 2021.

Pairwise DiD

I will go through the same 3 methods we used to study cannabis use, starting with the pairwise difference-in-differences. The number of possible graphs is much higher in this section since each of the 9 legal states has outcomes in 8 different drug classes to look at. I will try to provide a representative selection of results in the paper, but the best way to explore the results is through [the online appendix](#).

The data on per Medicaid enrollee prescription rates are noisier than the cannabis consumption data, so it is difficult to find a parallel trends match for some states. Data on less common medications, like spasticity and seizures, are noisier than for depression or pain medication. For almost all the states and drug classes, the observed rate of prescription medicine consumption is lower than the counterfactual predicted by an appropriately parallel control. Most of the results are not statistically significant. There are a few exceptions, like Colorado for depression and sleep meds, but we would expect a few statistically significant results from a collection this large even if the true effect was zero. Taken together, the preponderance of same-direction results and the smattering of significant ones suggest a mild but still detectable impact of cannabis legalization, and thus consumption, on prescription medicine habits. In all the graphs below, the independent variable is the number of units of medication (usually pills) prescribed per Medicaid enrollee.

Figure 9: Colorado and Utah Depression Pills Per Enrollee TWFE

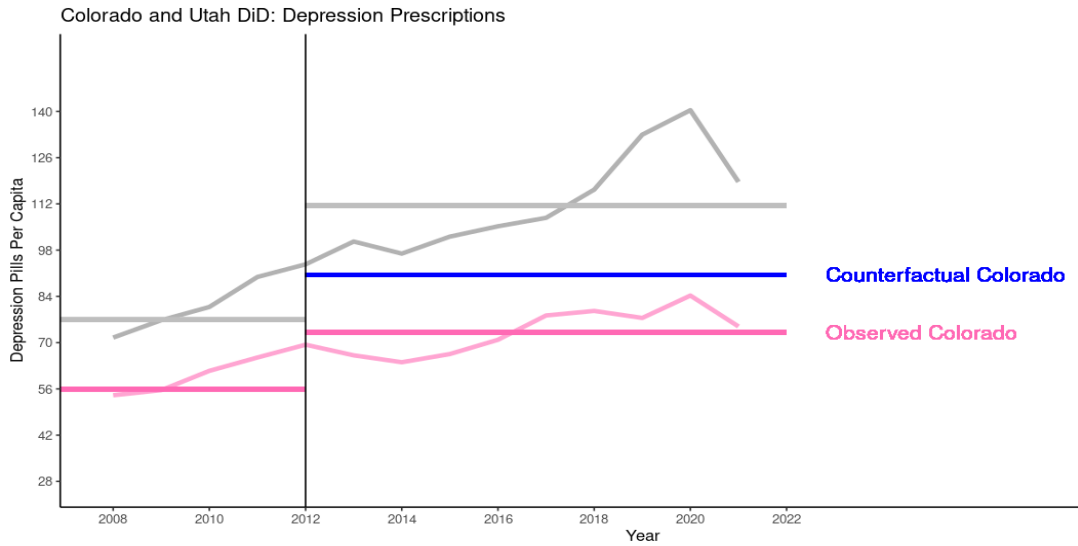


Table 3: Colorado and Utah Depression Pills Per Enrollee TWFE

Dependent variable: <i>PerCapUnits_reimbursed</i>				
	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	76.923	4.326	17.781	< 2e-16***
treated	-20.943	6.118	-3.423	0.00192**
time	34.539	5.472	6.312	7.94e-07***
treated:time	-17.372	7.739	-2.245	0.03286*

Note: Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
 Residual standard error: 10.6 on 28 degrees of freedom
 Multiple R-squared: 0.813, Adjusted R-squared: 0.793
 F-statistic: 40.57 on 3 and 28 DF, p-value: 2.52e-10

Figure 10: Oregon and Minnesota Sleep Pills Per Enrollee TWFE

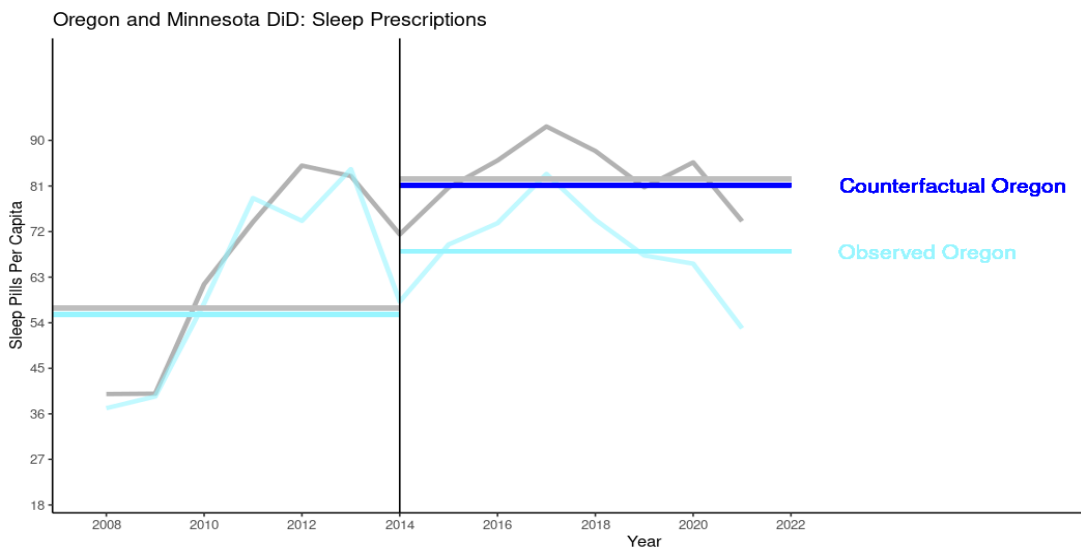


Table 4: Oregon and Minnesota Sleep Pills Per Enrollee TWFE

Dependent variable: <i>PerCapUnits_reimbursed</i>				
	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	56.892	5.693	9.994	9.76e-11***
treated	-1.232	8.051	-0.153	0.87947
time	25.532	8.051	3.171	0.00366**
treated:time	-13.086	11.385	-1.149	0.26014

*Note: Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1*
Residual standard error: 16.1 on 28 degrees of freedom
Multiple R-squared: 0.3383, **Adjusted R-squared:** 0.2674
F-statistic: 4.771 on 3 and 28 DF, **p-value:** 0.008263

Synthetic Controls

These pairwise DiD results have the same problems we faced in the first section: an informal process for selecting control states and only a vague sense of how all the results aggregate.

Finding an appropriate control state is sometimes obvious when there is a close match, but with some state-drug pairs it can be impossible to find another state which is parallel. Pairwise DiD is essentially the same as synthetic controls except the weighted average of control states is restricted to a vector of all zeros and a single one. If we drop this restriction we can greatly expand the space of prescription medicine trend lines that we can match.

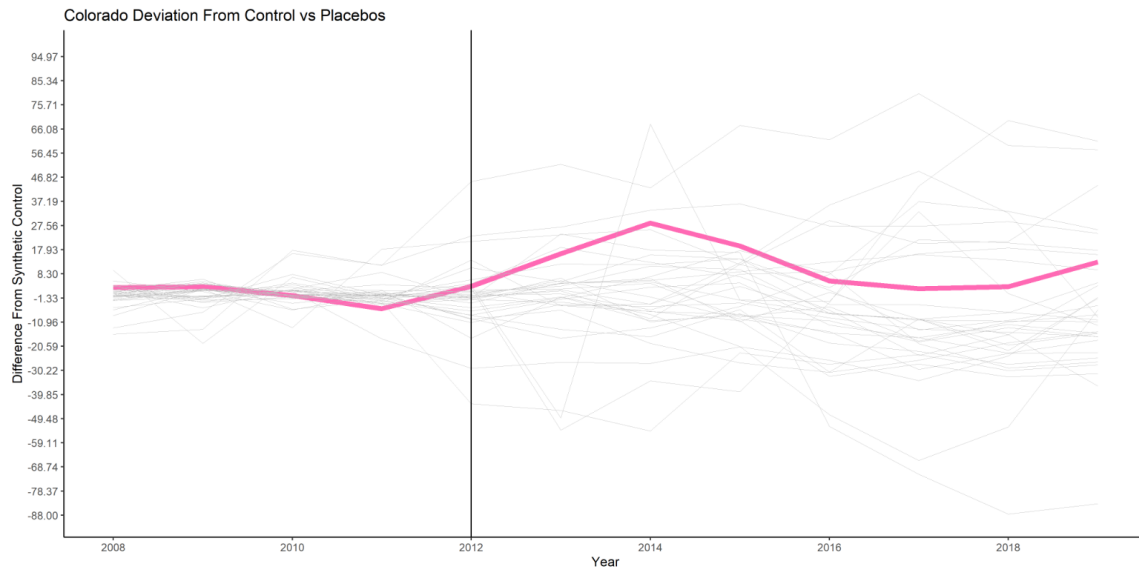
The results from synthetic control are milder than they were for cannabis consumption. There is no divergence between treatment and control as stark as Colorado or Nevada in cannabis consumption. When there is a large divergence, it is always in the direction predicted by a substitution relationship: consumption of prescription meds decreases in the legal state relative to its non-legal control. There is also some inconsistency with the pairwise DiD results. For

example, Colorado saw a significant decrease in depression medicine consumption relative to Utah in a pairwise DiD but it remains well matched by its synthetic control after legalization.

Figure 11: Colorado Depression Pills Per Enrollee Synthetic Control



Figure 12: Colorado Depression Pills Per Enrollee Compared to Placebos



The drug classes that see the largest and most consistent decreases across states are sleep and nausea.

Figure 13: Washington Nausea Pills Per Enrollee Synthetic Control

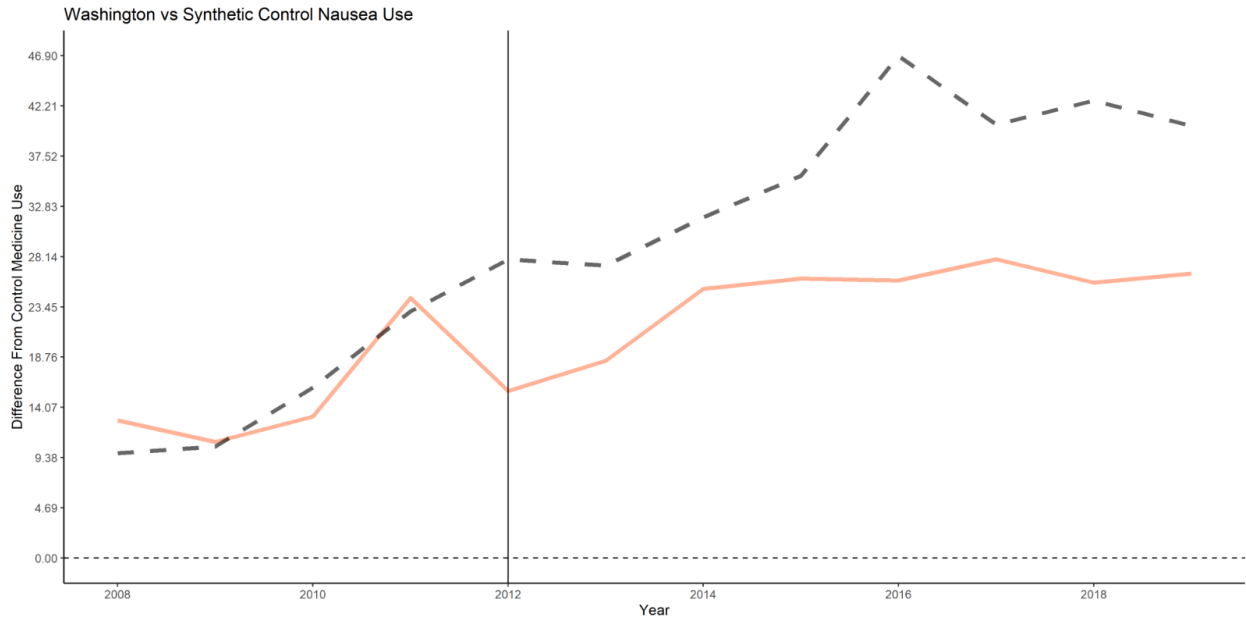
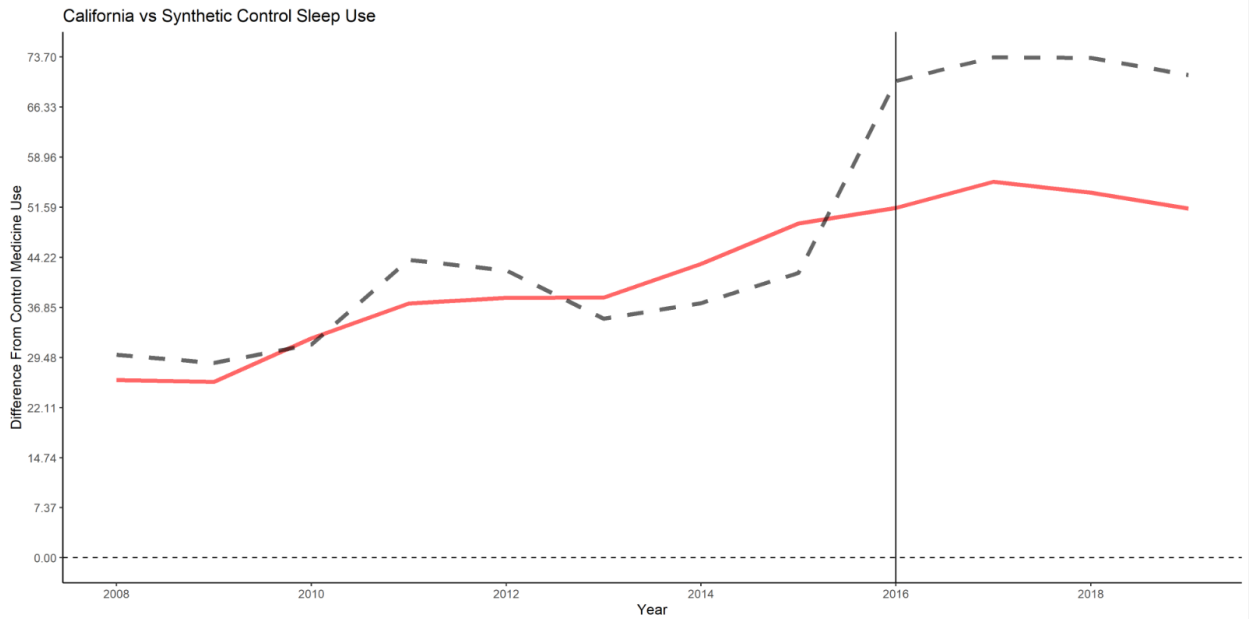
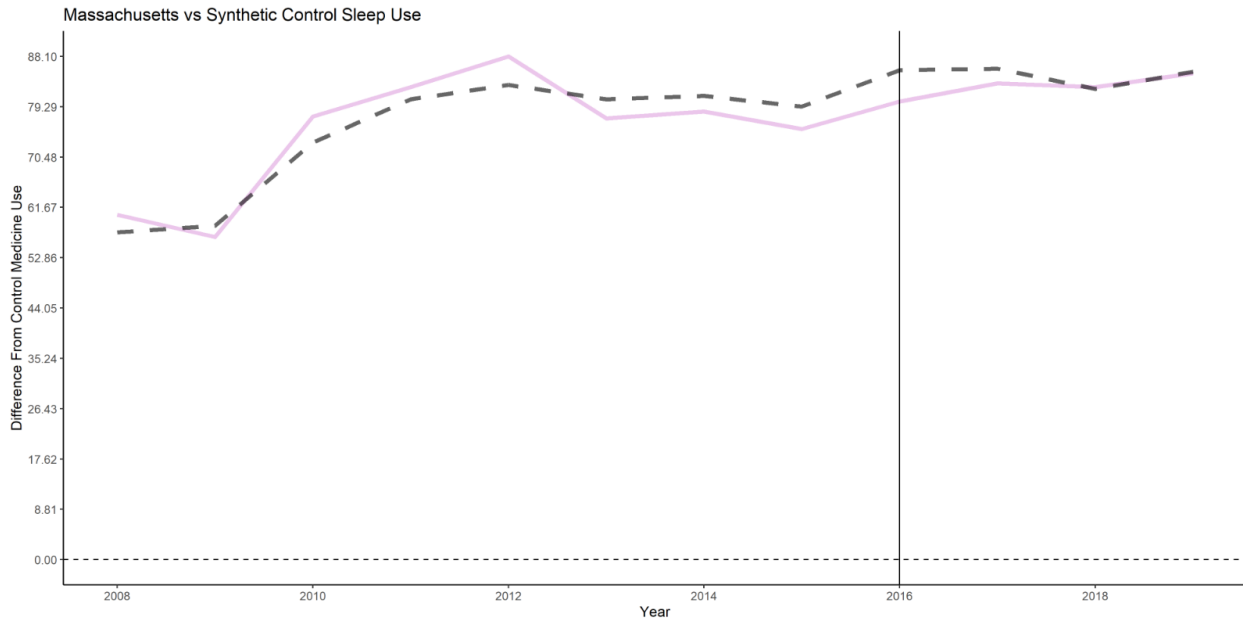


Figure 14: California Sleep Pills Per Enrollee Synthetic Control



But even these results do not show up in all of the states.

Figure 15: Massachusetts Sleep Pills Per Enrollee Synthetic Control



The synthetic control results seem smaller and more inconsistent than the pairwise DiDs. This should decrease our confidence in there being a causal impact of cannabis consumption on prescription meds but it isn't enough to conclude with confidence either way.

Event Study

We still need a way to aggregate all the results across states. A statistically insignificant result does not give zero information. If we have 9 insignificant results across all the legal states but their point estimates all show a decrease in medicine consumption it is much more likely that the true effect is negative than is shown by the p-value on any of the individual states' estimates. If the true effect was centered around zero it would be strange to have such consistent directional agreement across several random draws. We can use the event study specification to aggregate the results with this in mind.

I find statistically significant decreases in medicine consumption for the depression and anxiety drug classes (which share many drugs), and pain. Sleep and psychosis are nearly significant as well. In all these cases the effect size 6 and 7 years out from the treatment year is the most negative and significant. This reflects the much larger effects for the group of states that legalized in 2014 (Alaska, Oregon, and Washington D.C) than for the 2016 or 2012 groups. Below I reproduce the event study graph for pain medication. As with all the graphs in this section, the effect is measured in the number of units (usually pills) prescribed per Medicaid enrollee.

Figure 16: Average Treatment Effect of Legalization on Pain Medicine Use




Table 5: Average Treatment Effect of Legalization on Pain Medicine Use

Overall summary of ATT's			
	ATT	Std. Error	[95% Conf. Int.]
	-18.5923	9.5602	[-37.33, 0.1453]
Dynamic Effects:			
Event time	Estimate	Std. Error	[95% Simult. Conf. Band]
-9	6.2970	6.4993	[-10.3365, 22.9305]
-8	-8.4320	4.7818	[-20.6699, 3.8058]
-7	-3.9158	5.1633	[-17.1300, 9.2985]
-6	-0.3653	7.4013	[-19.3071, 18.5766]
-5	-0.5966	5.2617	[-14.0626, 12.8695]
-4	-3.7387	8.6465	[-25.8675, 18.3901]
-3	6.2118	9.3977	[-17.8393, 30.2629]
-2	-7.9544	15.8037	[-48.4003, 32.4916]
-1	6.2093	11.8176	[-24.0349, 36.4536]
0	-17.9993	11.6305	[-47.7648, 11.7662]
1	-18.7631	7.6943	[-38.4548, 0.9286]
2	-16.1440	9.7629	[-41.1298, 8.8417]
3	-11.8349	13.1753	[-45.5540, 21.8842]
4	-15.8228	15.7596	[-56.1558, 24.5101]
5	-15.5312	13.4647	[-49.9910, 18.9285]
6	-38.3320	11.9540	[-68.9253, -7.7386]*
7	-35.0571	11.3419	[-64.0839, -6.0303]*
8	-11.1578	12.3421	[-42.7446, 20.4289]
9	-5.2808	13.4789	[-39.7769, 29.2153]

Control Group: Never Treated, Anticipation Periods: 0

Estimation Method: Doubly Robust

Signif. codes: '*' confidence band does not cover 0

See [the online appendix for all of the results](#) and the event study in logs to estimate the effect in terms of a percentage change in medicine consumption.

Using our estimate of how much cannabis use increased after legalization and the estimated treatment effect (in percentage terms) of legalization on prescription medicines, we can estimate the elasticity of substitution between cannabis and prescription medications. According to the logged version of the event study, depression and pain medications saw about a 10% decrease in the number of pills prescribed per Medicaid enrollee and anxiety had a 9% decrease. Since cannabis use increased by 30% on average in these states, our best guess at the elasticity of

substitution between cannabis and these classes of prescription medicines is negative 1/3rd. A one percent increase in cannabis use is associated with a .33% reduction in prescription medicine use.

Discussion

Comparison With Previous Research

I find smaller effects than most previous research. The paper that is nearest in data sources and methods is this 2022 paper by Shyam Raman and Ashley Bradford [Raman, Bradford, \(2022\)](#).

We use the same classification of drugs and study the same population of Medicaid enrollees.

They find statistically significant reductions for 6 out of the 8 drug classes I study and estimate larger reductions in consumption of about 10%.

Given the similarity of our data sources I am not sure what explains the discrepancy in our results. There are a few factors which are probably important. Raman and Bradford only use one method, the event study, to get their results. They also only use the standard event study specification rather than the more robust version recommended by Callaway and Sant'Anna.

There are states which never get treated in the dataset, however, so the biases of this standard specification are not as severe. Second, they only include 3 years of data post treatment.


Finally, 'Non-Standard Errors' add significant variation to estimates on the same data between different teams of economists [Menkveld, et al. \(2021\)](#). This 2021 paper with hundreds of authors gives 164 teams of economists the same dataset and hypotheses to test. They found that differences in the way research teams got to their estimates created as much variation as the randomness of the data generating process itself. So the standard errors reported by each estimate

are about half as large as they would be if you took the uncertainty from the research process itself into account. Small differences in decisions made during data cleaning, collection, and analysis accumulate into sometimes large differences in results. This may explain part of the discrepancy between my paper and Raman and Bradford's piece.

My paper is also relevant to the ongoing research surrounding cannabis and opioids. Several papers have posited substitution relationships between cannabis and opioids using data both on opioid prescription [Liang, et al. \(2018\)](#) and opioid overdose [Bachhuber, et al. \(2014\)](#). However, a 2019 paper [Shover, et al. \(2019\)](#) attempted a replication and extension of the opioid overdose paper and found a positive and statistically significant effect, completely reversing the earlier finding. Both papers essentially use a version of the standard event study model to track changes in opioid overdoses following medical cannabis laws. My findings fall between these two with a mild negative effect of cannabis legalization on pain medication.

Publication bias may be an issue as well. A data rich question like this one has probably been investigated by hundreds of researchers. But due to the incentives and preferences of researchers,

Figure 17: Distribution of z-values from Medline 1976-2019 (Zwer, Cator, 2020)



there is a selection effect which biases the results that a literature reviewer can see. Researchers who find small or insignificant effects are not motivated to pursue the question further. Personal preferences for more interesting results and outside incentives for statistical significance work together to draw these researchers to other projects. Researchers who find large effects in the expected direction or complete reversals of previous work have discovered something interesting and publishable [Zwer, Cator \(2020\)](#). This selection for results on the tails of the true distribution of effects compounds with selection for results on the tails of the ‘Non-Standard Error’ distribution that arises from differences in data cleaning, evidence selection, choice of time window, and many other choices made at researcher’s discretion. Thus even if the true effect is mild, published literature reflects an outsized sample from the farthest edges of possible effect size estimates.

Previous studies have used fewer tests and less rigorous methods than this paper so my results should increase confidence in mild substitution effects between cannabis and prescription medications. However, placing high overall credence on my estimates is not the correct response. Just because my results are mild does not mean I am immune to the vagaries of random variation in both data generation and my own research methods. The primary takeaway ought not to be the movement of our best guess at the true effect towards my estimates but rather greater uncertainty.

Limitations

There are several limitations of this paper. First is the lack of distinction between intensive and extensive growth in cannabis use. The measure I use is from a survey by the National Survey on Drug Use and Health which asks “have you used cannabis in the last month?” The percentage of

respondents answering yes certainly increased following legalization but this is only one dimension of increased cannabis use. The question does not distinguish between someone who uses cannabis once a month and a daily user. If cannabis legalization increased the amount that existing users smoked it would not be captured by this measure. This poses difficulties for interpreting the ‘elasticity of substitution’ I calculated in the previous section. My measure of prescription drug use is an intensive measure: the number of pills prescribed per Medicaid enrollee.

Another limitation is that substitution effects for Medicaid enrollees may not generalize to the rest of the population. Medicaid enrollees differ from the general population in several important characteristics including age, race, income, and employment status. It is likely that both medical and recreational drug use interact with these characteristics in unique ways. Perhaps wealthier increase cannabis consumption at the expense of savings rather than trading off with medical expenses which they see as more essential. Or Medicaid users’ cannabis consumption may be more or less responsive to legal changes than the general population.

An assumption of all of the quasi-experimental methods that I use is “no anticipation.” This assumption is satisfied only if the treated group cannot see their treatment coming. If they do, then their strategic adaptations to the treatment can bias the results. This is relevant to my context since legislative changes like cannabis legalization may be anticipated, announced in advance, or unofficially rolled out early as law enforcement chooses not to pursue cannabis use. The event studies had almost all of the pre-treatment effect estimates centered on zero which suggests that there was not much anticipation, but it is still important to point out.

Policy Implications

On the national level, the most important policy implication of this paper is that the uncertainty in existing research should be acknowledged before important policy decisions are made based on the relationship between cannabis and prescription medicine use. Given the results of my research and some backing from intuition and economic theory, it would not be surprising if there truly was some substitution between cannabis and certain prescription medications.

However, the restriction of my data to Medicaid enrollees and the wide variation of results found in papers studying similar questions should temper confidence in extensions to national level policy.

Possible impacts on state budgets, both through substitution away from Medicaid services and from cannabis tax revenue, are more within the purview of my data. There are no statistically significant cost savings for Medicaid across all drug classes or for any drug class individually, although all of them have negative point estimates. Taking the point estimate of about 10\$ saved per Medicaid enrollee and, perhaps naively, multiplying it across all 84 million Medicaid enrollees we get a cost savings estimate of about \$840 million which is similar to the estimate of \$1.01 billion from David Bradford's 2016 paper [Bradford \(2017\)](#).

On the tax revenue side, the large and statistically significant increases in cannabis use that we estimated in the first section suggest big gains to government budgets. This 2021 paper from the National Tax Journal [Miller, Seo \(2021\)](#) finds that while there are overall tax revenue increases, cannabis' substitution with other taxed goods, especially alcohol and cigarettes, cuts the increase by 40%.

Conclusion

Cannabis legalization was a widespread and rapid shift in drug policy in the US. The full implications of this shift are not well understood. It is clear that legalization significantly increases cannabis use by at least 30% and this effect may continue to increase across the nation to the higher levels seen in Colorado and Washington. As more states legalize and national legalization becomes more likely, second order effects of this increase in cannabis use are important to consider. Previous research has found that cannabis has substitution effects with a number of prescription medications. My research extends the evidence on these substitution effects and increases confidence in their existence, although it tempers previous estimates of their magnitude. Substitution with prescription drugs could be a compelling reason to legalize cannabis nationally due to cost savings in medical care, increased tax revenue, and easier access to palliative care. All of these effects are consistent with the point estimates in my paper. However, the confidence intervals are wide enough that running experiments and collecting more data while implementing legalization is worth doing.

References

- Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2), 391–425.
<https://doi.org/10.1257/jel.20191450>
- Bachhuber, M. A., Saloner, B., Cunningham, C. O., & Barry, C. L. (2014). Medical cannabis laws and opioid analgesic overdose mortality in the United States, 1999-2010. *JAMA Internal Medicine*, 174(10), 1668. <https://doi.org/10.1001/jamainternmed.2014.4005>
- Blessing, E. M., Steenkamp, M. M., Manzanares, J., & Marmar, C. R. (2015). Cannabidiol as a potential treatment for anxiety disorders. *Neurotherapeutics*, 12(4), 825–836.
<https://doi.org/10.1007/s13311-015-0387-1>
- Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.
<https://doi.org/10.1016/j.jeconom.2020.12.001>
- Feingold, D., & Weinstein, A. (2020). Cannabis and depression. *Cannabinoids and Neuropsychiatric Disorders*, 67–80. https://doi.org/10.1007/978-3-030-57369-0_5
- Liang, D., Bao, Y., Wallace, M., Grant, I., & Shi, Y. (2018). Medical cannabis legalization and opioid prescriptions: Evidence on US Medicaid enrollees during 1993-2014. *Addiction*, 113(11), 2060–2070. <https://doi.org/10.1111/add.14382>
- Lucas, P., & Walsh, Z. (2017). Medical Cannabis Access, use, and substitution for prescription opioids and other substances: A survey of authorized medical cannabis

patients. *International Journal of Drug Policy*, 42, 30–35.

<https://doi.org/10.1016/j.drugpo.2017.01.011>

Menkveld, A. J. (2021). Non-standard errors. *SSRN Electronic Journal*.

<https://doi.org/10.2139/ssrn.3979358>

Miller, K., & Seo, B. (2021). The effect of cannabis legalization on substance demand and tax revenues. *National Tax Journal*, 74(1), 107–145. <https://doi.org/10.1086/712915>

Powell, D., Pacula, R. L., & Jacobson, M. (2018). Do medical marijuana laws reduce addictions and deaths related to pain killers? *Journal of Health Economics*, 58, 29–42.

<https://doi.org/10.1016/j.jhealeco.2017.12.007>

Powell, D., Pacula, R. L., & Jacobson, M. (2018). Do medical marijuana laws reduce addictions and deaths related to pain killers? *Journal of Health Economics*, 58, 29–42.

<https://doi.org/10.1016/j.jhealeco.2017.12.007>

Raman, S., & Bradford, A. C. (2022). Recreational cannabis legalizations associated with reductions in prescription drug utilization among Medicaid enrollees. *Health Economics*, 31(7), 1513–1521. <https://doi.org/10.1002/hec.4519>

Shover, C. L., Davis, C. S., Gordon, S. C., & Humphreys, K. (2019). Association between medical cannabis laws and opioid overdose mortality has reversed over time. *Proceedings of the National Academy of Sciences*, 116(26), 12624–12626.

<https://doi.org/10.1073/pnas.1903434116>

Subbaraman, M. S. (2016). Substitution and complementarity of alcohol and cannabis: A review of the literature. *Substance Use & Misuse*, *51*(11), 1399–1414.

<https://doi.org/10.3109/10826084.2016.1170145>

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, *225*(2), 175–199.

<https://doi.org/10.1016/j.jeconom.2020.09.006>

Whiting, P. F., Wolff, R. F., Deshpande, S., Di Nisio, M., Duffy, S., Hernandez, A. V., Keurentjes, J. C., Lang, S., Misso, K., Ryder, S., Schmidtkofer, S., Westwood, M., & Kleijnen, J. (2015). Cannabinoids for medical use. *JAMA*, *313*(24), 2456.

<https://doi.org/10.1001/jama.2015.6358>

Zwet, E. W., & Cator, E. A. (2021). The significance filter, the winner's curse and the need to shrink. *Statistica Neerlandica*, *75*(4), 437–452. <https://doi.org/10.1111/stan.12241>